PORTAL
THE ACM DIGITAL LIBRARY

> home > about > feedback > login

US Patent & Trademark Office

Try the *new* Portal design
Give us your opinion after using it.

Search Results

Search Results for: **[sql <paragraph> random]**
Found **81** of **124,998 searched.**

09/864,591

Search within Results

GO > Advanced Search

> Search Help/Tips

Sort by: **Title** **Publication** **Publication Date** **Score** Binder

Results 1 - 20 of 81 **short listing**

Prev Page 1 2 3 4 5 Next Page

1 Security for Web Applications and P2P: Abstracting application-level web security 100%
David Scott , Richard Sharp
Proceedings of the eleventh international conference on World Wide Web May 2002
Application-level web security refers to vulnerabilities inherent in the code of a web-application itself (irrespective of the technologies in which it is implemented or the security of the web-server/back-end database on which it is built). In the last few months application-level vulnerabilities have been exploited with serious consequences: hackers have tricked e-commerce sites into shipping goods for no charge, user-names and passwords have been harvested and condential information (such as ...

2 Developing a valuable tool for IT staff: increasing efficiency and communication with an intranet 100%
Ryan LeBreton , Lisa Veloz
Proceedings of the 31st annual ACM SIGUCCS conference on User services September 2003
During the summer of 2002, Information Services & Resources (ISR) created a departmental web team with the goal of redesigning the ISR website. The team completed a successful (and SIGUCCS award winning) redesign, but in the process discovered a need for an organized website for the IT staff. At the same time our department's organizational development team was looking for a way to distribute their materials to staff over the web and build an online community. We decided to combine efforts and d ...

3 Another breadth-first approach to CS I using python 100%
Christine Shannon


ACM SIGCSE Bulletin , Proceedings of the 34th SIGCSE technical symposium on Computer science education January 2003
Volume 35 Issue 1

In an effort to serve the needs of both majors and non-majors, the Computer Science Department at Centre College has restructured the CS I course so that it uses the language Python, devotes more attention to the Internet and the World Wide Web, addresses ethical and societal issues, and introduces students to programmable robots and an SQL database. This diverse course has been attractive to the students while still maintaining a strong emphasis on programming.

4 A benchmark of NonStop SQL release 2 demonstrating near-linear speedup and scaleup on large databases 100%


Susanne Englert , Jim Gray , Terrye Kocher , Praful Shah
ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1990 ACM SIGMETRICS conference on Measurement and modeling of computer systems April 1990
Volume 18 Issue 1

5 Expert design tools for physical database design 100%


Rajiv Tewari
Proceedings of the 1990 ACM SIGBDP conference on Trends and directions in expert systems September 1990

6 Consistancy and replication: TiNA: a scheme for temporal coherency-aware in-network aggregation 100%


Mohamed A. Sharaf , Jonathan Beaver , Alexandros Labrinidis , Panos K. Chrysanthis
Proceedings of the 3rd ACM international workshop on Data engineering for wireless and mobile access September 2003

This paper presents TiNA, a scheme for minimizing energy consumption in sensor networks by exploiting end-user tolerance to temporal coherency. TiNA utilizes temporal coherency tolerances to both reduce the amount of information transmitted by individual nodes (communication cost dominates power usage in sensor networks), and to improve quality of data when not all sensor readings can be propagated up the network within a given time constraint. TiNA was evaluated against a traditional in-network ...

7 Invited talk: Analyzing massive data streams: past, present, and future 100%


Minos Garofalakis
Proceedings of the 8th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery June 2003

Continuous data streams arise naturally, for example, in the installations of large telecom and Internet service providers where detailed usage information (Call-Detail-Records, SNMP-/RMON packet-flow data, etc.) from different parts of the underlying network needs to be continuously collected and analyzed for interesting trends. Such environments raise a critical need for effective stream-processing algorithms that can provide (typically, approximate) answers to data-analysis queries while util ...

8 XML query processing I: Dynamic sample selection for approximate query processing 100%


Brian Babcock , Surajit Chaudhuri , Gautam Das
Proceedings of the 2003 ACM SIGMOD international conference on on Management of data June 2003

In decision support applications, the ability to provide fast approximate answers to aggregation queries is desirable. One commonly-used technique for approximate query answering is sampling. For many aggregation queries, appropriately constructed biased (non-uniform) samples can provide more accurate approximations than a uniform sample. The optimal type of bias, however, varies from query to query. In this paper, we describe an approximate query processing technique that dynamically constructs ...

9 Query processing: Estimating compilation time of a query optimizer 100%



Ihab F. Ilyas , Jun Rao , Guy Lohman , Dengfeng Gao , Eileen Lin

Proceedings of the 2003 ACM SIGMOD international conference on on Management of data June 2003

A query optimizer compares alternative plans in its search space to find the best plan for a given query. Depending on the search space and the enumeration algorithm, optimizers vary in their compilation time and the quality of the execution plan they can generate. This paper describes a compilation time estimator that provides a quantified estimate of the optimizer compilation time for a given query. Such an estimator is useful for automatically choosing the right level of optimization in comme ...

10 Online analytic processing (OLAP): Spreadsheets in RDBMS for OLAP 100%



Andrew Witkowski , Srikanth Bellamkonda , Tolga Bozkaya , Gregory Dorman , Nathan Folkert , Abhinav Gupta , Lei Shen , Sankar Subramanian

Proceedings of the 2003 ACM SIGMOD international conference on on Management of data June 2003

One of the critical deficiencies of SQL is lack of support for n-dimensional array-based computations which are frequent in OLAP environments. Relational OLAP (ROLAP) applications have to emulate them using joins, recently introduced SQL Window Functions [18] and complex and inefficient CASE expressions. The designated place in SQL for specifying calculations is the SELECT clause, which is extremely limiting and forces the user to generate queries using nested views, subqueries and complex joins ...

11 Perspective: Improving the efficiency of database-system teaching 100%



Jeffrey D. Ullman

Proceedings of the 2003 ACM SIGMOD international conference on on Management of data June 2003

The education industry has a very poor record of productivity gains. In this brief article, I outline some of the ways the teaching of a college course in database systems could be made more efficient, and staff time used more productively. These ideas carry over to other programming-oriented courses, and many of them apply to any academic subject whatsoever. After proposing a number of things that could be done, I concentrate here on a system under development, called OTC (On-line Testing Cente ...

12 Measurement: Automatically inferring patterns of resource consumption in network traffic 100%



Cristian Estan , Stefan Savage , George Varghese

Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications August 2003

The Internet service model emphasizes flexibility -- any node can send any type of traffic at any time. While this design has allowed new applications and usage models to flourish, it also makes the job of network management significantly more

challenging. This paper describes a new method of traffic characterization that automatically groups traffic into minimal clusters of conspicuous consumption. Rather than providing a static analysis specialized to capture flows, applications, or network-to ...

13 A parallel pipelined relational query processor: An architectural overview 100%

Daniel Gajski , Won Kim , Shinya Fushimi

ACM SIGARCH Computer Architecture News , Proceedings of the 11th annual international symposium on Computer architecture January 1984

Volume 12 Issue 3

This paper outlines the overall architecture of a query processor for relational queries and describes the design and control of its major processing modules. The query processor consists of only four processing modules and a number of random-access memory modules. Each processing module processes tuples of relations in a bit-serial, tuple-parallel manner for each of the primitive database operations which comprise a complex relational query. The query processor is designed to be manufactur ...

14 A new approach to protein structure and function analysis using semi-structured databases 100%

William M. Shui , Raymond K. Wong , Stephen C. Graham , Lawrence K. Lee , W. Bret Church

Proceedings of the First Asia-Pacific bioinformatics conference on Bioinformatics 2003 - Volume 19 January 2003

The development of high-throughput genome sequencing and protein structure determination techniques have provided researchers with a wealth of biological data. Integrated analysis of such data is difficult due to the disparate nature of the repositories used to store this biological data and of the software used for its analysis. This paper presents a framework based upon the use of semi-structured database management systems that would provide an integrated interface for the collection, storage ...

15 A Pareto model for OLAP view size estimation 100%

Thomas P. Nadeau , Toby J. Teorey

Proceedings of the 2001 conference of the Centre for Advanced Studies on Collaborative research November 2001

On Line Analytical Processing (OLAP) aims at gaining useful information quickly from large amounts of data residing in a data warehouse. To improve the quickness of response to queries, pre-aggregation is a useful strategy. However, it is usually impossible to pre-aggregate along all combinations of the dimensions. The multi-dimensional aspects of the data lead to combinatorial explosion in the number and potential storage size of the aggregates. We must selectively pre-aggregate. Cost/benefit a ...

16 An SQL interface to X.500 100%

David Barrowman , Patrick Martin

Proceedings of the 1995 conference of the Centre for Advanced Studies on Collaborative research November 1995

The X.500 standard specifies a distributed directory service designed to store information about people and objects associated with computer networks. Its API is geared toward retrieving information based on this application domain. Recently, a number of projects have used the directory in non-traditional ways. Such

applications, however, are constrained by the X.500 information model and the limited functionality of its API. We describe a prototype system that allows users to view the informati ...

17 Internet intrusions: global characteristics and prevalence 100%



Vinod Yegneswaran , Paul Barford , Johannes Ullrich
ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2003 ACM SIGMETRICS international conference on Measurement and modeling of computer systems June 2003
Volume 31 Issue 1

Network intrusions have been a fact of life in the Internet for many years. However, as is the case with many other types of Internet-wide phenomena, gaining insight into the *global* characteristics of intrusions is challenging. In this paper we address this problem by systematically analyzing a set of firewall logs collected over four months from over 1600 different networks world wide. The first part of our study is a general analysis focused on the issues of distribution, categorization ...

18 Establishing the semantic web 1: SemTag and seeker: bootstrapping the semantic web via automated semantic annotation 100%



Stephen Dill , Nadav Eiron , David Gibson , Daniel Gruhl , R. Guha , Anant Jhingran , Tapas Kanungo , Sridhar Rajagopalan , Andrew Tomkins , John A. Tomlin , Jason Y. Zien
Proceedings of the twelfth international conference on World Wide Web May 2003

This paper describes Seeker, a platform for large-scale text analytics, and SemTag, an application written on the platform to perform automated semantic tagging of large corpora. We apply SemTag to a collection of approximately 264 million web pages, and generate approximately 434 million automatically disambiguated semantic tags, published to the web as a label bureau providing metadata regarding the 434 million annotations. To our knowledge, this is the largest scale semantic tagging effort to ...

19 Information retrieval 2: Text joins in an RDBMS for web data integration 100%



Luis Gravano , Panagiotis G. Ipeirotis , Nick Koudas , Divesh Srivastava
Proceedings of the twelfth international conference on World Wide Web May 2003

The integration of data produced and collected across autonomous, heterogeneous web services is an increasingly important and challenging problem. Due to the lack of global identifiers, the same entity (e.g., a product) might have different textual representations across databases. Textual data is also often noisy because of transcription errors, incomplete information, and lack of standard formats. A fundamental task during data integration is matching of strings that refer to the same entity. ...

20 Astrolabe: A robust and scalable technology for distributed system monitoring, management, and data mining 100%



Robbert Van Renesse , Kenneth P. Birman , Werner Vogels
ACM Transactions on Computer Systems (TOCS) May 2003
Volume 21 Issue 2

Scalable management and self-organizational capabilities are emerging as central requirements for a generation of large-scale, highly dynamic, distributed applications. We have developed an entirely new distributed information management system called Astrolabe. Astrolabe collects large-scale system state, permitting rapid updates

and providing on-the-fly attribute aggregation. This latter capability permits an application to locate a resource, and also offers a scalable way to track sys ...

Results 1 - 20 of 81 **short listing**

Prev Page 1 2 3 4 5 Next Page

The ACM Portal is published by the Association for Computing Machinery. Copyright ? 2003 ACM, Inc.